arXiv:1702.06383v1 [cs.IR] 21 Feb 2017

# Deep Geometric Retrieval

Y. Qian*[1], E Vazquez[1], and B Sengupta†[1,2]

[1]Cortexica Vision Systems Limited, London SE1 8RT, UK
[2]Dept. of Engineering, University of Cambridge, Cambridge CB2 1PZ, UK

**Abstract**

Comparing images in order to recommend items from an image-inventory is a subject of continued interest. Added with the scalability of deep-learning architectures the once 'manual' job of hand-crafting features have been largely alleviated, and images can be compared according to features generated from a deep convolutional neural network. In this paper, we compare distance metrics (and divergences) to rank features generated from a neural network, for content-based image retrieval. Specifically, after modelling individual images using approximations of mixture models or sparse covariance estimators we resort to their information-theoretic and Riemann geometric comparisons. We show that using approximations of mixture models enable us to to compute a distance measure based on the Wasserstein metric that requires less effort than computationally intensive optimal transport plans; finally, an affine invariant metric is used to compare the optimal transport metric to its Riemann geometric counterpart – we conclude that although expensive, retrieval metric based on Wasserstein geometry are more suitable than information theoretic comparison of images. In short, we combine GPU scalability in learning deep feature vectors with computationally efficient metrics that we foresee being utilized in a commercial setting.

## 1 Introduction

A common problem in computer vision lies in finding similarity between 2 (or 3)-dimensional images (or tensors). This is pursed by measuring distances between the two objects, primarily using normalized co-relation, Euclidean distance, Bhattacharyya distance, Jensen-Shannon divergence, amongst many others. The distances are measured after the images are encoded in some latent space wherein such a latent structure is learnt using a variety of classifiers – support vector machines (SVMs), logistic regression, etc. Recently, due to the advantages of scalability, large-scale classifier frameworks based on deep-learning have been used for music recommendation [28], image recommendation [23] as

*yu.qian@cortexica.com; BS and YQ contributed equally to this paper.
†bs573@cam.ac.uk; BS has a dual appointment at Dept. of Bioengineering, Imperial College London.

well as general recommendation architectures [5]. Most of these framework do not take the underlying geometry of the feature space into account while making recommendation. This becomes increasingly important when similarity between objects is measured in terms of 'perception', a quantity that is oblivious to the commonly used distance metrics. The non-trivial problem lies in (a) collecting perceptual similarity between objects in a database (via psychophysics), and (b) using this similarity to construct a metric for classification and retrieval. To compound the problem further, metric used for comparing images might be very different than those for comparing sounds. Yet, regardless of the modality, a large stream of work in neuroscience hypothesize that perception is based on minimizing the prediction error between what is actually observed and what we predict we will observe [11].

In this paper, we start with the $0^{th}$ order problem i.e., compare how different distance metrics fair against one another when objects that are to be compared are represented as probabilistic objects in the brain; furthermore, we use approximations of the probability density to compute a metric based on the principle of optimal transport [29] and Riemann geometry [1] that takes in account the geometry of transport between two images, an idea that is inherently important for solving the perceptual similarity problem [24].

Technically, the problem lies in searching an image database (ranking) with millions of image features ($A_{n_i \times r_m \times f_j}$) for sets of images (say up to 10-20 images) that have similar properties to a query image ($\hat{a}_{query}$). Here, $A$ is a $n_i \times r_m \times f_j$ tensor where $n_i$ is the total number of images, $r_m$ is the length of each feature vector and $f_j$ is the total number of features extracted from one image; $\hat{a}_{query}$ is a $r_m \times f_j$ query matrix. Although, one can manually construct feature-vectors based on wavelet decomposition, low-rank approximations, etc., we rely on using a convolution neural network (CNN) to compute the feature signature ($r_m \times f_j$) for images in the database, as well as the query.

One way to operationalize a solution lies in weighting each image in the database using a weight vector, and subsequently extremize the mutual information (or another comparison metric) between the query and the database with respect to the weights. The result of this optimization problem leads us to a weight vector that provides a rank for all the images in the database when compared to the query image. This is equivalent to measuring distances where each image lies on a continuous probability manifold. There are two contribution of this paper – (a) in order to describe each image with its deep feature, we use either a computationally efficient approximation of Gaussian Mixture Models (GMMs) or a sparse covariance estimator based on Given rotations, and (b) we provide comparison between these probability distributions using a variety of information-theoretic and geometric metrics. This work leads us to a much deeper problem where geometric similarity measures can be possibly combined to approximate the metric governing **perceptual similarity**.

# 2 Methods

## 2.1 Dataset and deep-feature generation

In this paper, Describable Textures Dataset (DTD) [6] is used to evaluate geometric similarity measures for image retrieval. Rather than recognition and

description of object, texture images in DTD are collected from wild images (Google and Flickr) and classified based on human visual perception [27], such as directionality (line-like), regularity (polka-dotted and chequered), etc. DTD is therefore selected in this research to evaluate the similarity measurements base on human visual perception. DTD contains 5640 wild texture images with 47 describable attributes drawn from the psychological literature and is publicly available on the web at `http://www.robots.ox.ac.uk/vgg/data/dtd/`.

Textures can be described via orderless pooling of filter bank response [12]. In Deep CNN, the convolutional layers are akin to non-linear filter banks; these have in fact been proved to be better for texture descriptions [7]. Here, the local deep features are extracted from last convolutional layer of a pre-trained VGG-M [4]. This is represented by $A = \left(a_1, ..., a_i, ..., a_N : a \in \mathbb{R}^D\right)$; the size of last convolutional layer is $H \times W \times D$, where $D$ denotes the dimension vector of filter response at the $i^{th}$ pixel of last convolution layer; $N = H \times W$ is the total number of local features. For image level representation, two methods are applied on local features – one is to generate a Gaussian Mixture Model (GMM) model on local descriptors and the second is to estimate a shrinked yet sparse co-variance matrix from the deep feature representation of individual images. A statistical similarity metric is then applied to rank images.

## 2.2 Retrieval and Ranking

In subsection 2.2.1-2.2.3, we describe three approaches to rank images in terms of their 'statistical similarity' (not perceptual similarity). For the first, we use an information-theoretic divergence whilst the second and third distances are based on the cost involved in transporting one image to another, and geodesic distance on a Riemannian manifold, respectively.

In order to rank images in the database we use two methods, one is to build a Gaussian Mixture Model (GMM) [20], and the second is to estimate a covariance matrix from deep features. For each image, we model the $r_m \times f_j$ feature matrix using a GMM. Specifically, for computational and analytical efficiency, we approximate the GMM with a Normal distribution, such that the sufficient statistics reads,

$$\begin{aligned} \tilde{\mu} &= \sum_a \omega_a \mu_a \\ \tilde{\Sigma} &= \sum_a \omega_a \left(\Sigma_a + (\mu_a - \tilde{\mu})(\mu_a - \tilde{\mu})^T\right) \end{aligned} \tag{1}$$

$\mu_a$, $\Sigma_a$ and $\omega_a$ are the mean, co-variance and the mixing weights of each Normal distributions (subscript $a$).

The second approximation to an image relies on estimating the co-variance matrix from the feature matrix generated from a deep convolutional neural network. Although, the geometry of the co-variance matrix can be utilized to estimate it using low-rank and sparse penalization, for the sake of computational efficiency, we use an alternative treatment due to Refs. [3, 2] i.e., a fast sparse matrix transformation (SMT). Briefly, the SMT imposes sparsity constraint on the manifold of co-variance matrices yet maintains a full-rank representation. This is useful as the computation is $\mathcal{O}(f_j)$; the SMT can also be seen as a generalization of FFT and orthonormal para-unitary wavelet transform.

We will assume that each feature vector is *i.i.d* zero mean Normal random vectors, and the sample covariance is simply, $\frac{1}{n}AA^T$; it is a unbiased estimate of the true covariance matrix, $R = \mathbb{E}[S] = E\Lambda E^T$. Often time $S$ is singular, and shrinkage estimators [15] are used to regularize the covariance matrix by shrinking it towards a target structure such as an identity matrix, a diagonal matrix with sample variances, amongst others. Sparsity can also be imposed, as in graphical lasso [10] by imposing a 1-norm constraint on the precision matrix. The maximum likelihood (ML) estimate of the eigenvectors ($E$) and the eigenvalues ($\Lambda$) gives us,

$$\begin{aligned} \hat{E} &= \underset{E \in \Omega_k}{\arg\min} \left\{ \left| diag\left(E^T S E\right) \right| \right\} \\ \hat{\Lambda} &= diag\left(\hat{E}^T S \hat{E}\right) \end{aligned} \tag{2}$$

The SMT constrains the feasible set of $\Omega_k$ to a set of orthonormal transformations that are selected as a SMT of order K. A matrix $E$ is a SMT of order K if it can be factorized to K sparse orthonormal matrices i.e.,

$$\begin{aligned} E &= \prod_{k=1}^{K} E_k = E_1 E_1 \dots E_k \\ E_k &= I + \Theta(i_k, i_j, \theta_k) \end{aligned} \tag{3}$$

Each sparse matrix $E$ can be constructed as a orthonormal Givens rotation on a pair of co-ordinate indexes $(i_k, i_j)$ of Givens rotations such that,

$$[\Theta]_{ij} = \begin{cases} \cos(\theta_k) - 1, \text{ if } i = j = i_k \text{ or } i = j = j_k \\ \sin(\theta_k), \text{ if } i = i_k \text{ and } j = j_k \\ -\sin(\theta_k), \text{ if } i = j_k \text{ and } j = i_k \\ 0, \text{ otherwise} \end{cases} \tag{4}$$

Using greedy minimization [3, 2] we have,

$$\begin{aligned} \hat{E}_k &= \arg\min \left| diag\left(E_k^T S_k E_k\right) \right| \\ S_{k+1} &= \hat{E}_k^T S_k \hat{E}_k \end{aligned}$$

$$\begin{aligned} \hat{E} &= \prod_{k=1}^{K} \hat{E}_k \\ \hat{\Lambda} &= diag\left(S_{k+1}\right) \end{aligned} \tag{5}$$

As a final step, we obtain a shrinked co-variance matrix where the shrinkage parameter $\alpha$ is selected using cross-validation,

$$\begin{aligned} \Sigma_{SMT} &= \hat{E}\hat{\Lambda}\hat{E}^T \\ \Sigma &= \alpha \cdot \Sigma_{SMT} + (1 - \alpha) \cdot S \end{aligned} \tag{6}$$

### 2.2.1 Ranking by KL-divergence

Since there is no analytical solution between the KL-divergence between two GMMs ($V_a \sim \mathcal{N}_a(\omega_a, \mu_a, \Sigma_a)$ and $(V_b \sim \mathcal{N}_b(\omega_b, \mu_b, \Sigma_b)$, we utilize two approximations: in the first we approximate the GMM with Eqn. 1. The KL-divergence ($D_{KL}\left(V_a \| V_b\right)$) now reads,

$$\frac{1}{2}\left[\log\frac{|\Sigma_b|}{|\Sigma_a|} - N_d + tr\left(\Sigma_b^{-1}\Sigma_a\right) + \left(\mu_b - \mu_a\right)^T\Sigma_b^{-1}\left(\mu_b - \mu_a\right)\right] \tag{7}$$

In our experiments, we compute a symmetric-KL divergence which is simply $D_{KL}^{Normal} = \frac{1}{2}D_{KL}\left(V_a \| V_b\right) + \frac{1}{2}D_{KL}\left(V_b \| V_a\right)$. Sorting the KL-divergence provides us with a similarity rank.

This is a gross-approximation wherein a more subtle approximation relies in bounding the KL-divergence. Particularly, using results from information theory [14, 21], we provide retrieval results using a variational approximation to the KL divergence. Particularly, since the log-function is concave, using Jensen's inequality we have,

$$\begin{aligned}
D_{KL}\left(V_a \parallel V_b\right) &= \mathbb{E}_{V_a}\left[V_a\right] - \mathbb{E}_{V_a}\left[V_b\right] \\
\mathbb{E}_{V_a}\left[V_b\right] &= V_a \log V_b \\
&= \sum_a \omega_a \int V_a \log \sum_b \phi_{b|a}\frac{\omega_b V_b}{\phi_{b|a}} \\
&\geqslant \sum_a \omega_a \int V_a \sum_b \phi_{b|a} \log\frac{\omega_b V_b}{\phi_{b|a}} \\
&= \sum_a \omega_a \sum_b \phi_{b|a}\left(\log\left(\frac{\omega_b}{\phi_{b|a}}\right) + \int V_a \log V_b\right)
\end{aligned} \tag{8}$$

Here, $\phi_{b|a}$ is a variational parameter that is positive and sums to one. Maximizing $w.r.t\ \phi_{b|a}$ yields,

$$\mathbb{E}_{V_a}\left[V_b\right] \geqslant \sum_a \omega_a \log \sum_b \omega_a e^{-D_{KL}\left(V_a \| V_b\right)} - \sum_a \omega_a \mathcal{H}\left(V_a\right) \tag{9}$$

$\mathcal{H}$ is the entropy functional. Subsequently, the variational bound becomes,

$$D_{KL}^{\text{variational}}\left(V_a \parallel V_b\right) = \sum_a \omega_a \log \frac{\sum_{a'} \omega_{a'} e^{-D_{KL}\left(V_a \| V_{a'}\right)}}{\sum_b \omega_b e^{-D_{KL}\left(V_a \| V_b\right)}} \tag{10}$$

We symmetrize the variational KL by using $D_{KL}^{\text{var}} = \frac{1}{2}D_{KL}^{\text{var}\,iational}\left(V_a \| V_b\right) + \frac{1}{2}D_{KL}^{\text{var}\,iational}\left(V_b \| V_a\right)$. Note that such a divergence is the difference of two variational approximations, not a bound in itself.

### 2.2.2 Ranking *via* Kantorovich relaxation

Let $(\Psi, \psi_m)$ and $(\Lambda, \lambda_m)$ denote two Polish probability spaces depicting *image 1* and *image 2*, respectively – $\psi_m$ and $\lambda_m$. The trivial coupling between the two exists if $\Psi$ and $\Lambda$ are independent so that the coupling is simply a tensor product $\psi_m \otimes \lambda_m$. A more useful coupling exists when there is a function $S : \Psi \to \Lambda$ such that $\lambda = S(\psi)$. The transport map $S$ is equivalently the change of variables from $\psi_m$ to $\lambda_m$.

**Definition of a transport map:** Let $S$ be a Borel map: $\Psi \to \Lambda$, the push forward of $\psi_m$ through $S$ is the Borel measure, denoted $S_{\#\psi_m}$ defined on $\Lambda$ by $S_{\#\psi_m}(\Lambda) = \psi_m(S^{-1}(\Lambda))$. A Borel map: $\Psi \to \Lambda$ is said to be a transport map if $S_{\#\psi_m} = \lambda_m$.

In optimal transport [29], there is a cost entailed by transporting one measure into another. The transport map then relies on finding the infimum of $\left(\int_\Psi c\left(x, S(x)\right) d\psi_m : S_{\#\psi_m} = \lambda_m\right)$. Optimal transference plans are important because such couplings are stable to perturbations, they encode geometric information about the underlying cost-function, and they exist in smooth as well as non-smooth settings. Given that the existence of this transport map can not be guaranteed, a Kantorovich relaxation amounts to a convex relaxation of Monge's formulation wherein we seek a coupling $\gamma \in P\left(\Psi, \Lambda\right)$,

$$\gamma_0 = \underset{\gamma \in P(\Psi,\Lambda)}{\arg\min} \int_{\Psi \times \Lambda} c\left(x^\psi, x^\lambda\right) d\gamma \tag{11}$$

The joint probability measure with the marginals $\psi_m$ and $\lambda_m$ allow us to define a Wasserstein distance of order $p$ between $\psi_m$ and $\lambda_m$,

$$\mathcal{W}\left(\psi_m, \lambda_m\right) = \inf\left(\left\{\mathbb{E}\ D\left(x^\psi, x^\lambda\right)^p\right\}^{1/p}\right) \tag{12}$$

$D$ is a distance with a corresponding cost of $c\left(x^\psi, x^\lambda\right) = d\left(x^\psi, x^\lambda\right)^p$. This Earth-Mover or the Monge-Kantorovich distance provides us with a metric over the space of squared integrable probability measures. For two Normal distribution, the $L_2$-Wasserstein distance [29, 26] reads,

$$D_{\mathcal{W}} = \|\mu_a - \mu_b\|_2^2 + tr\left(\Sigma_a + \Sigma_b - 2\left(\Sigma_a^{1/2}\Sigma_b\Sigma_a^{1/2}\right)^{1/2}\right) \tag{13}$$

Once the GMMs have been approximated via Eqn. 1 or Eqn. 6, it is fairly simple to compute distances using Eqn. 13.

### 2.2.3 Ranking *via* Affine Invariant Riemannian Metric

Let us again consider two feature matrices (query and database), $V_a \sim \mathcal{N}(0, \Sigma_a)$ and $V_b \sim \mathcal{N}(0, \Sigma_b)$. These positive-definite matrices are elements of $S_{++}^{f \times f}$, a space with a defined Riemannian metric [9, 19]. Under such a geometry, the distance $D_R\left(V_a, V_b\right)$ between these two matrices is,

$$D_R(\Sigma_a, \Sigma_b) = \left\| \log(\Sigma_a^{-1/2} \Sigma_b \Sigma_a^{-1/2}) \right\|_F = \left[ \sum_{c=1}^{C} \log^2 \lambda_c \right]^{1/2} \tag{14}$$

$c$ is the dimension of the co-variance matrix, $\lambda_c$ are the eigen-values and $F$ represents the Frobenius norm. A useful property of such a distance is that regardless of how the images are manipulated – be it re-scaling, normalization, whitening, filtering, etc. – the distance between the two sources as captured by Eqn. 14 remains invariant.

To compute Eqn. 14 one can use Eqn. 1 to approximate both GMMs as Normal distributions; alternatively, the covariance estimated using Eqn. 6 can be used.

# 3 Experiments

In this section, deep feature geometric retrieval methods are evaluated on the DTD dataset. For each image, a set of local deep features is extracted from last convolutional layer of a pretrained VGG-M. The dimension of each local feature vector is 512. A GMM with 64 components is subsequently generated from the set of local deep features. Normal approximation by GMM and sparse covariance estimation by SMT are used to represent the feature matrix; information-theoretic (Normal and Variational approximation KL) and geometric (Wasserstein and Riemannian) measures gauge the similarity of two images. In this experiment, Normal approximation KL, Variational approximation KL and Wasserstein metric are applied on GMM model respectively and represented by GMM-Normal KL, GMM-Variational KL and GMM-Wasserstein. Normal approximation KL, Wasserstein and Remannian metric are applied on sparse covariance generated by SMT respectively and denoted by SMT-Normal KL, SMT-Wasserstein and SMT-Riemannian.

| MAP | Top-1 | Top-5 | Top-10 | Time |
|---|---|---|---|---|
| GMM-Normal KL | 0.53 | 0.46 | 0.42 | 0.375s |
| GMM-Variational KL | 0.45 | 0.42 | 0.38 | 0.016s |
| GMM-Wasserstein | 0.62 | 0.52 | 0.46 | 5.147s |
| SMT-Riemannian | 0.50 | 0.44 | 0.39 | 0.754s |
| SMT-Normal KL | 0.53 | 0.44 | 0.39 | 0.125s |
| SMT-Wasserstein | 0.59 | 0.51 | 0.46 | 9.04s |

Table 1: Retrieval results on the DTD dataset

To evaluate similarity metric for image retrieval, Mean Average Precision (MAP) on top 10 rankings are calculated. 2 images per category i.e., a total of 94 images are selected as queries from the test dataset. The dataset retrieved includes 3760 images from DTD training and validation datasets. MAP on DTD is listed in Table 1. Average precision on top 10 ranking is displayed in Figure 1. An example of the retrieval obtained with each method is shown in Figure 2. On each case 10 images are displayed. Top left image is the query used. The

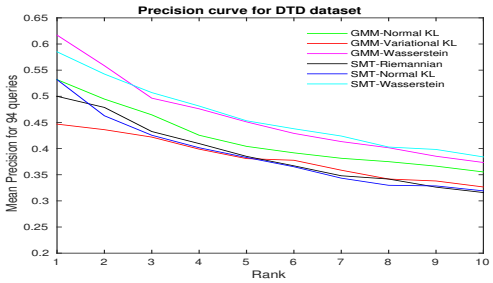


Figure 1: Precision on the DTD dataset

Figure 2: Retrieved results on DTD: a) GMM-Wasserstein b) SMT-Wasserstein c) GMM-Normal KL d) SMT-Normal KL e) SMT-Riemannian f) GMM-Variational KL

rest of images are arranged by similarity to query image as obtained with each method.

Query times shown in Tabel 1 have been obtained as the average time to calculate similarity between two images. The code was implemented in Matlab 2015a under linux with Intel(R) Xeon(R) CPU E5-2640 @ 2.00GHz and 125G RAM.

# 4 Conclusion

In this paper, we have touched upon the $0^{th}$ order problem that leads to understanding 'perceptual similarity'. More specifically, we have used a convolution neural network (CNN) to obtain feature matrices; utilizing either Gaussian Mixture Models (GMMs) or shrinked covariance estimators to obtain a probabilistic representation of the features. Subsequently, using information theoretic divergences and geometric metrics (see Appendix A) we compare (dis) similarities between images. Based on evaluation for DTD dataset, Wasserstein distances show increased retrieval fidelity based on both probabilistic representations, albeit they are more expensive to evaluate. We believe that the increased fidelity of the Wasserstein distance is due to two properties – first, the metric does not include calculating the inverse of covariance matrices, thereby enclosing the cases with singularity; in contrast, the KL-divergence between two distributions could easily reach infinity if the covariance of the second distribution becomes singular. The second property, which we hypothesize, is the increased statistical

robustness of the Wasserstein distance i.e., the metric might have small variance when comparing distributions that are closely situated in the parametric manifold.

Although, we have utilized the final convolutional layer of a CNN to distinguish images; a lot of empirical work has shown that there are many general features of an image or a video that are captured by the initial layer of a CNN [30]. By visualizing different layers in Ref. [18], it is apparent that the lower layer of CNN can capture more colour information, the higher layers on the other hand are more objective. The retrieval result in Figure 2 demonstrate that colour is not adequately captured due to local deep features extracted from the last convolutional layer, which keeps less colour related information. Hence, the fidelity to distinguish images using any of our retrieval criteria should undoubtedly increase with additional 'independent' feature vector that can be computed via the initial or the middle (general to more specific characterization of the image) layers of a CNN.

**Factors that affect the successful deployment**

For a commercial system, speed is an essential ingredient. In fact, computing Wasserstein and Riemann distance have their own issues. For example, Wasserstein distance in computer vision has been proposed more than a decade ago [22]. The cost of computing optimal transport between two distributions of dimension $d$ is at least $\mathcal{O}(d^3 \log d)$. This is especially not plausible to compute in a commercial environment when features vectors are generated by deep convolution neural networks, which are by construction high dimensional. In our study, even after approximating the GMMs as multivariate Normal distributions, the computational inefficiency is inherent, as computing Eqn. 13 proves to be most expensive amongst all the metrics that we compare. A solution emerges in the form of low dimensional embedding of the metric space [13, 16]; such solutions introduce distortions in addition to an increase in computational cost when the embedding dimension becomes larger than 4 [8]. Additionally, they are not designed to be scalable to take advantages of large-scale GPU resources. Ref. [8] has suggested to improve the scalability of the distance calculation by using an iterative diagonal scaling algorithm, known as Sinkhorn's algorithm or alternatively iterative proportional fitting. We leave this scalability issue for future work.

Similarly, computing the geodesic distance between two co-variance matrices is equally time inefficient – $\mathcal{O}(4d^3)$. The main component of this inefficiency emerges from the generalized eigenvalue equation, particularly for calculating multiple Cholesky factorization each time a query is initiated. One way forward may be to use Stein's distance [25] while preserving affine invariance and geometric properties inherited by the covariance matrices. Another way ahead is to perform the factorization on a GPU [17]. This becomes increasingly important if our framework were to be used for indexing similar videos (instead of images). This future application relies on returning a set of similar videos in response to a query video. This could replace the current text based tagged video framework, as that used by several online video platforms, with feature based tagged videos.

# References


[1] S. Amari, H. Nagaoka, and D. Harada. *Methods of Information Geometry*. Translations of Mathematical Monographs. American Mathematical Society, 2007.

[2] Guangzhi Cao, Leonardo R. Bachega, and Charles A. Bouman. The sparse matrix transform for covariance estimation and analysis of high dimensional signals. *IEEE Trans. Image Processing*, 20(3):625–640, 2011.

[3] Guangzhi Cao, Charles A. Bouman, and Kevin J. Webb. Noniterative MAP reconstruction using sparse matrix representations. *IEEE Trans. Image Processing*, 18(9):2085–2099, 2009.

[4] K. Chatfield, K. Simonyan, A. Vedaldi, and A. Zisserman. Return of the devil in the details: Delving deep into convolutional nets. In *British Machine Vision Conference*, 2014.

[5] Heng-Tze Cheng, Levent Koc, Jeremiah Harmsen, Vihan Jain, Xiaobing Liu, and Hemal Shah. Wide & deep learning for recommender systems. In *Proceedings of the 1st Workshop on Deep Learning for Recommender Systems*, Deep Learning for Recommender Systems 2016, pages 7–10, New York, NY, USA, 2016. ACM.

[6] M. Cimpoi, S. Maji, I. Kokkinos, S. Mohamed, and A. Vedaldi. Describing textures in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2014.

[7] M Cimpoi, S Maji, and A Vedaldi. Deep filter banks for texture recognition and segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2015.

[8] Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in Neural Information Processing Systems*, pages 2292–2300, 2013.

[9] Wolfgang Förstner and Boudewijn Moonen. *A Metric for Covariance Matrices*, pages 299–309. Springer Berlin Heidelberg, Berlin, Heidelberg, 2003.

[10] Jerome Friedman, Trevor Hastie, and Robert Tibshirani. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441, 2008.

[11] Karl Friston, Biswa Sengupta, and Gennaro Auletta. Cognitive dynamics: From attractors to active inference. *Proceedings of the IEEE*, 102(4):427–445, 2014.

[12] Yunchao Gong, Liwei Wang, Ruiqi Guo, and Svetlana Lazebnik. Multi-scale orderless pooling of deep convolutional activation features. In *Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part VII*, pages 392–407, 2014.

[13] Kristen Grauman and Trevor Darrell. Fast contour matching using approximate earth mover's distance. In *Proceedings of the IEEE Computer Society Conference on CVPR*, volume 1, 2004.

[14] John R Hershey and Peder A Olsen. Approximating the Kullback Leibler divergence between Gaussian mixture models. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, volume 4, 2007.

[15] W. James and James Stein. Estimation with quadratic loss. In Jerzy Neyman, editor, *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability*, pages 361–379, 1961.

[16] Haibin Ling and Kazunori Okada. An efficient earth mover's distance algorithm for robust histogram comparison. *IEEE transactions on Pattern Analysis and Machine Intelligence*, 29(5):840–853, 2007.

[17] Gary Macindoe. *Hybrid algorithms for efficient Cholesky decomposition and matrix inverse using multicore CPUs with GPU accelerators*. PhD thesis, University College London, 2013.

[18] A. Mahendran and A. Vedaldi. Understanding deep image representations by inverting them. In *Proceedings of the IEEE Computer Society Conference on CVPR*, 2015.

[19] M. Moakher. A differential geometric approach to the geometric mean of symmetric positive-definite matrices. *SIAM J. Matrix Anal. Appl.*, 26:35–747, 2005.

[20] Kevin P Murphy. *Machine learning: a probabilistic perspective*. MIT Press, Cambridge, MA, 2012.

[21] F. Nielsen and K. Sun. Guaranteed bounds on the Kullback-Leibler divergence of univariate mixtures. *IEEE Signal Processing Letters*, 23(11):1543–1546, Nov 2016.

[22] Yossi Rubner, Carlo Tomasi, and Leonidas J Guibas. A metric for distributions with applications to image databases. In *Computer Vision, 1998. Sixth International Conference on*, pages 59–66. IEEE, 1998.

[23] B Sengupta, E Vazquez, V Simaiaki, M Sasdelli, Y Qian, M Peniak, L Netherton, and G Delfino. Cortexica – a scalable image analysis framework. (under review), 2017.

[24] Biswa Sengupta, Arturo Tozzi, Gerald K Cooray, Pamela K Douglas, and Karl J Friston. Towards a neuronal gauge theory. *PLoS Biol*, 14(3):e1002400, 2016.

[25] S. Sra. Positive definite matrices and the s-divergence, 2015.

[26] Asuka Takatsu. On Wasserstein geometry of the space of Gaussian measures. *arXiv preprint arXiv:0801.2250*, 2008.

[27] H. Tamura, S. Mori, and T. Yamawak. Textural features corresponding to visual perception. *IEEE Transactions on systems, Man and Cybernetics*, page 460 473, 1978.

[28] Aaron Van den Oord, Sander Dieleman, and Benjamin Schrauwen. Deep content-based music recommendation. In *Advances in Neural Information Processing Systems*, pages 2643–2651, 2013.

[29] Cedric Villani. *Optimal Transport: Old and New*. Springer-Verlag, 2009.

[30] Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. How transferable are features in deep neural networks? In *Advances in NIPS 27*, pages 3320–3328. 2014.

# A Appendix

**Definition of a metric:** If $\mathcal{D}$ characterizes the distance between two objects it is a metric *iff* the following axioms hold: (a) $\mathcal{D}(a,b) \geqslant 0$, (b) $\mathcal{D}(a,a) = 0$, (c) $\mathcal{D}(a,b) = 0$ if $a = b$, (d) $\mathcal{D}(a,b) = \mathcal{D}(b,a)$ and (e) $\mathcal{D}(a,c) \leqslant \mathcal{D}(a,b) + \mathcal{D}(b,c)$. Semi-metrics only adhere to (a-b) and (d).